Introduction

Introduction to fMRI

Functional magnetic resonance imaging (fMRI) is an imaging technique that can be used to view various parts of the brain in real time. Physicians use fMRI in clinical settings as a non-invasive, radiation-free technique to image normal or abnormal brain activity. It has also been used to provide valuable insight to psychologists and neuroscientists, who wish to study neuronal functions as a patient or participant performs various activities.

In order to measure neural activity in the brain, fMRI measures changes in blood flow through a technique known as Blood-oxygen-level dependent (BOLD) contrast. Because the brain does not store glucose, whenever neurons in a particular area activate, glucose must be delivered at a greater rate to the site to supply sufficient energy for this stimulation. In order to do so, blood flow in this region of the brain increases. BOLD contrast measures this change in blood flow and allows us to map neuronal activity indirectly by understanding the associated changes in flow rates. However, the changes in blood flow lag the increase in neuronal activity and must be adjusted for when considering brain activity in real time.

Problem Statement

In this experiment, we aim to use fMRI to determine how localization of activity in the brain changes as our mentor Ricky Savjani watches different types of videos, which have been separated into 6 functional categories: action, animals, buildings, landscapes, objects, and people. Through analysis of the BOLD-contrast imaging, we seek to find local areas of stimulation of distinct portions of the brain that may result from watching these different types of videos.

Background

Physical chemistry basis

Functional magnetic resonance imaging (fMRI) takes advantage of the chemical principle of nuclear magnetic resonance to produce data which map activated brain regions over time. Atoms with an odd number of protons or neutrons have an intrinsic magnetic moment [1]. When exposed to a magnetic field, the atoms will emit electromagnetic radiation with a frequency proportional to the strength of the magnetic field until the magnetic poles align with the field [2]. Varying this magnetic field linearly along one direction will produce linearly varying frequencies. The position of the atom in this dimension is proportional to the frequency of the signal emitted. Radio frequency pulses along slices in one direction provide further spatial detail in the x and y direction [1]. At each pulse the atoms are temporally shifted to align with the altered magnetic field direction. The time it takes for the atom to realign with the stable magnetic field produce signals of varying frequencies [1].

Image Reconstruction with Fourier Transforms

The signal acquired from the electromagnetic emission of atoms contains information about the magnitude of each volume element, or voxel. A Fourier transform is performed on the time signals gathered from the emitted electromagnetic radiation to transfer it to the frequency domain [3]. The signal is dismantled and organized in the frequency domain. The frequency is related to its position based on how the magnetic field and radio frequency pulses were applied. An inverse Fourier transform transforms the signals from the frequency domain to 3D space, creating the actual fMRI image [3]. In this space, each unit of resolution is a 3D voxel and the magnitude of each voxel corresponds to the amplitude of the response at that region.

Brain mapping

A map of brain activity can be formed by taking MRI scans over time (i.e. fMRI). Blood oxygen levels are the quantities of interest in an fMRI scan. Blood with higher oxygen concentration has a different frequency response than blood with lower oxygen levels [2]. Active regions of the brain will quickly receive oxygen-rich blood. This spike in blood flow is called the

hemodynamic response, a plot of which is shown in [link]. An fMRI scan can determine where these active regions are over time and thus map the active regions of the brain for particular tasks. This type of imaging is called blood-oxygen-level dependent (BOLD) contrast imaging [2]. The result of an fMRI scan is a 4D matrix containing the intensity of each voxel. The first three dimensions correspond to space and the fourth is the time dimensions. In essence, this provides pictures of the brain over time; higher values in a location indicate that that region of the brain is active.

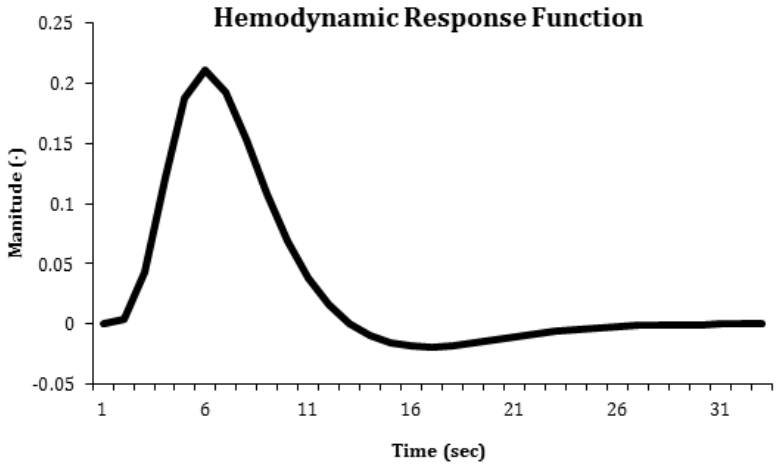

The hemodynamic response function (HRF).

Approach

To conclude whether a region of the brain had been activated by a video, regression analysis was performed on the fMRI data to determine which regions of the brains were involved in eliciting a certain response pattern. If a movie clip activates a region of the brain, that voxel value should switch from a low to a high value, essentially turning “on”. Fitting parameters from a regression indicate whether this “on” response occurred for any voxel for any movie.

Experimental Setup

While in an MRI scanner, a subject watched 6 types of videos on average 10 seconds long. The types were assigned based on what the main subject matter of the video was and were as follows: object, animal, people, action, landscape, and building. The back portion of the brain was scanned every second. This region contains the visual cortex, which is stimulated when watching video. The region of the scan is the blue region displayed in [link]. A coronal slice of this region is displayed in [link].

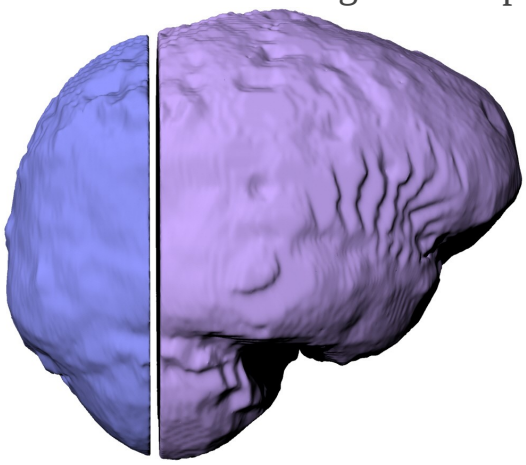

3D image of the subject's brain. Scanned region in blue.

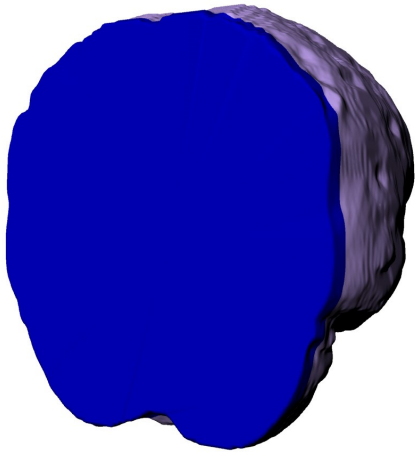

Blue slice shows coronal view of the brain.

Generalized Linear Model

The GLM is a means to identify significant parameters that affect an experimental response. Parameters are made to fit the measured data through multiplication with beta values. The formulation is as follows:

$$\beta$$

In this formulation Y represents some measured quantity over time, X is the predictor of what that value will be and β is a fitting parameter chosen to minimize the error between X and Y. In general, X can have several columns representing different predictors of the output and each predictor would then have a corresponding β value. In this analysis only one predictor is implemented. The β value signifies the weight that a certain predictor has in producing the output Y.

This approach can be applied to fMRI data to determine what regions “turn on” in response to a particular type of movie. The response of a given voxel during a certain movie is the measured response, Y. The predictor is the expected response of a voxel that has just been activated by a new video clip. In the simplest case, this would be modeled by a unit step – a strictly on/off response. The β values are parameters that indicate whether or not a

certain voxel exhibits this “on” response. A high β value for a given voxel shows that that voxel has been activated for that movie clip.

Convolution

The actual value of a voxel over time correlates to the blood flow in that region of the brain. Blood flow cannot start and stop instantaneously, as modeled by the unit step. The actual blood response to activation in the brain is best modeled by the hemodynamic response function (HRF). Convolution is a mathematical operation that can be used to find the output of a system given an input and the impulse response function of the system. Treating the HRF as the impulse response function to the system and the unit step as the input to the system, the output can be calculated through convolution [link]. Convolving these two signals results in a more biologically accurate predictor for the measured fMRI data. The expected response from the convolution, shown below in [link], is inserted as the predictor in the GLM instead of the unit step in order to arrive at more accurate β values.

β values were calculated for each voxel for 15 movies of each type. The higher the beta value for a voxel, the more weight that voxel has in creating the expected response. Arranging the β values in 3 dimensions yields a map of the active regions of the brain during a movie clip.

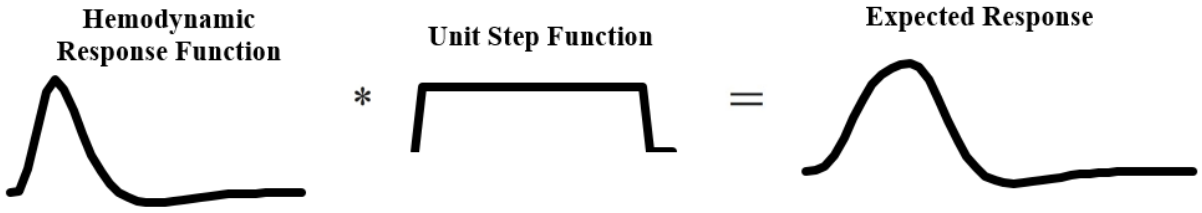

The convolution of the HRF and unit step to yield a better predictor, X, for the GLM.

The Algorithm

Three computational steps accomplished the goal of creating maps of active regions of the brain for each type of movie. First the HRF was convolved with a normalized unit step function. The result was truncated to match the

duration of each of the video clips. Next, β values were calculated for each voxel for 19 movies of each type. For one movie clip each voxel received one β value, indicating its activation during that clip. β values were averaged for each of the six movie types. Finally, these vectors of β values were rearranged into 3D space to create a map of brain activation. The full process produces 6 brain activity maps corresponding to 6 types of video clips.

Results

Discussion

Images collected from the preprocessed fMRI were combined into a 64x64x18x5783 matrix, where each point represents an (x,y,z,t) value for the BOLD response. Each voxel can therefore be represented by an (x,y,z) component and analyzed across the time domain to determine the response to each video type. To do so, the matrix was split up into different, smaller matrices of size 64x64x18x(duration) by the video and the video type. The first 15 videos of each type were analyzed.

For each video, each of the 73728 voxels was fit to the convolution of the unit step and the hemodynamic response function over the duration of the video using a generalized linear model (GLM). This produced fit parameters β for each voxel for each video which were then averaged according to video type and plotted for multiple cross sections. These cross sections were then weaved together into videos as shown below in [link]. Recall that only a section of the brain was imaged in this case to improve spatial and temporal resolution

From top to bottom, these images show the activity of the brain in (a) action, (b) animal, (c) building, (d) landscape, (e) object, and (f) people videos. Positive correlations are shown in red, while negative correlations are shown in blue. Videos were formed by moving from the relative front of the brain to the back. Recall that only the rear portion of the brain is shown.

Action Videos

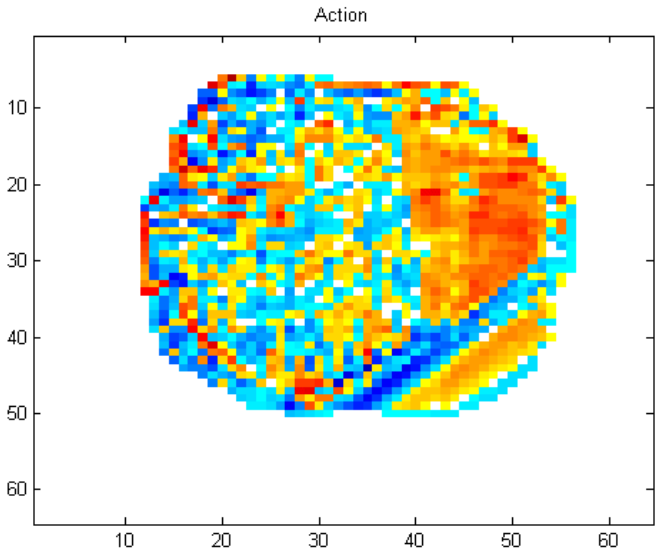

Animal Videos

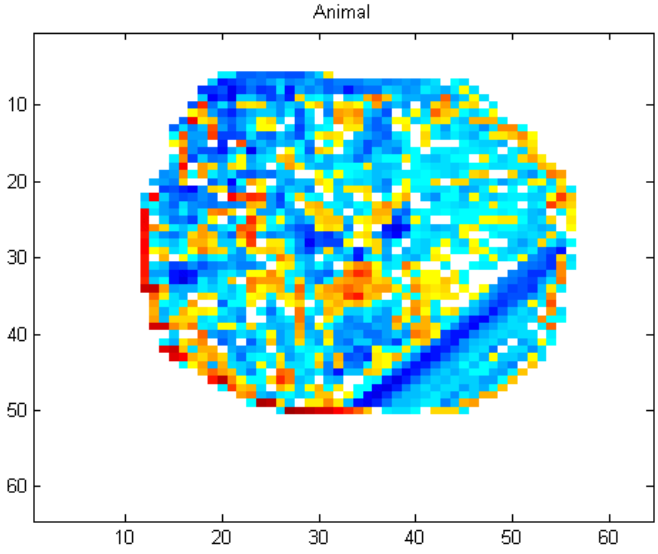

Building Videos

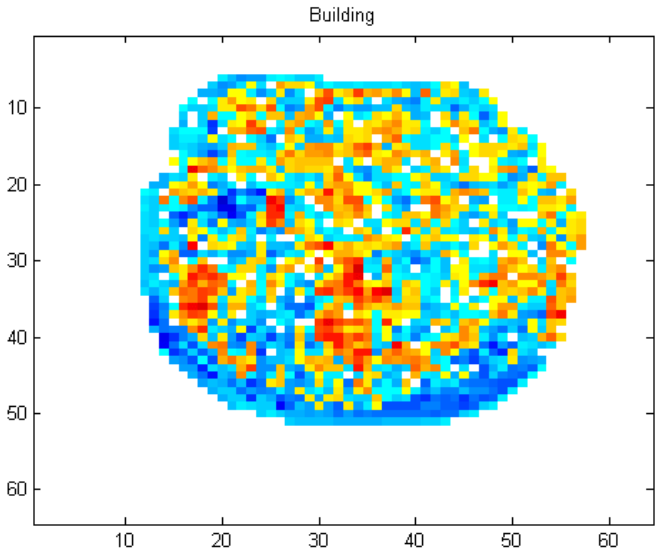

Landscape Videos

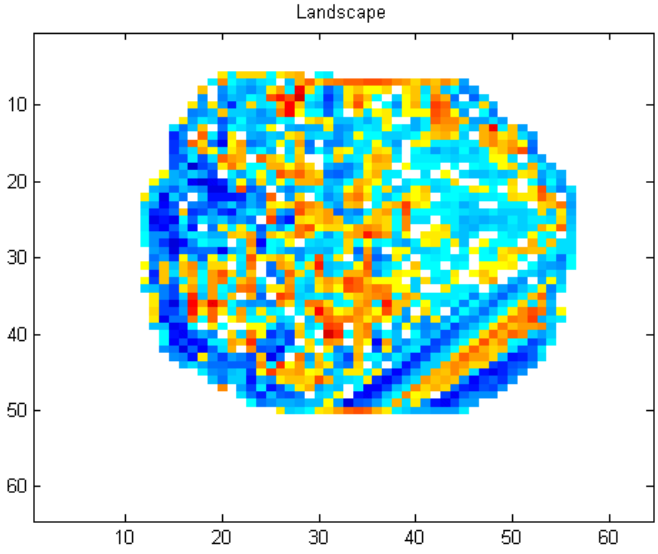

Object Videos

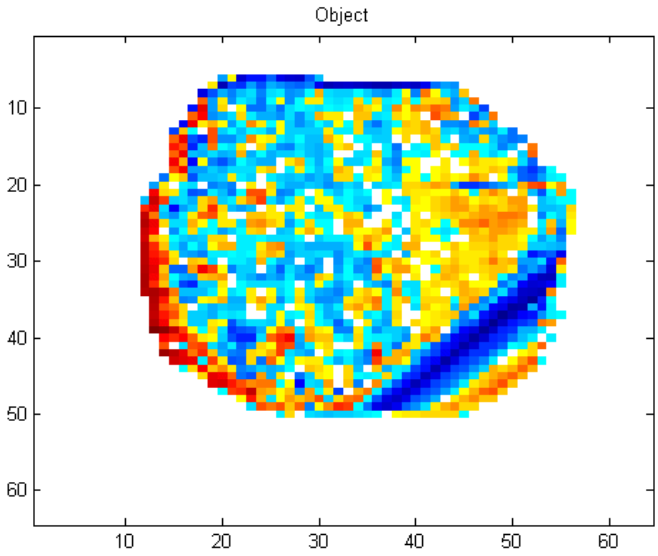

People Videos

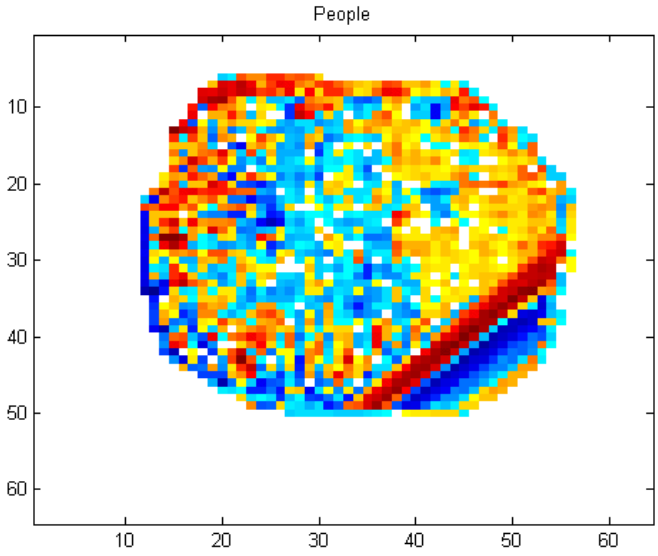

From inspection, we can see that there are areas of heightened stimulation that differ across video types. In order to confirm this analytically, we calculated an ANOVA table for the means of the 6 different videos and determined that the stimulation in response to the different types is significantly different ( ).

As validation of our analysis, we noticed in every video a relatively high level (red) of stimulation of the parietal and occipital lobes, which are areas of visual processing within the brain. We expect this to occur, since the

visual stimulation should activate the areas of the brain responsible for image processing and analysis within the brain.

Some interesting characteristics that we noted from visual inspection was a heightened stimulation in every section of the action videos when compared to the other videos, which could indicate the high level of integration of all of the other types of videos and the level of processing that must occur to capture the details within these videos. Large zones of little to no stimulation are also exhibited in the building, landscape and animal videos, which may indicate decreased interest during these videos. However, note that most of this analysis is qualitative in nature, and may not accurately reflect future analysis of the correlations.

Conclusions

Discussion

Through analysis of fMRI data we are able to distinguish the localized stimulation of different types of videos. By determining the value of the correlation between the convolution of the hemodynamic response function and the unit step function with the BOLD signal at each voxel, we can effectively determine the relative response of each voxel to these videos.

Challenges

We noticed that our analysis runs into issues with the edges of the brain as shown below in [link]. The lines and large flat regions of high/low correlation are uncharacteristic of this type of data and disappears upon moving one voxel into the image. This may result from inaccuracies in the preprocessing techniques at the edges.

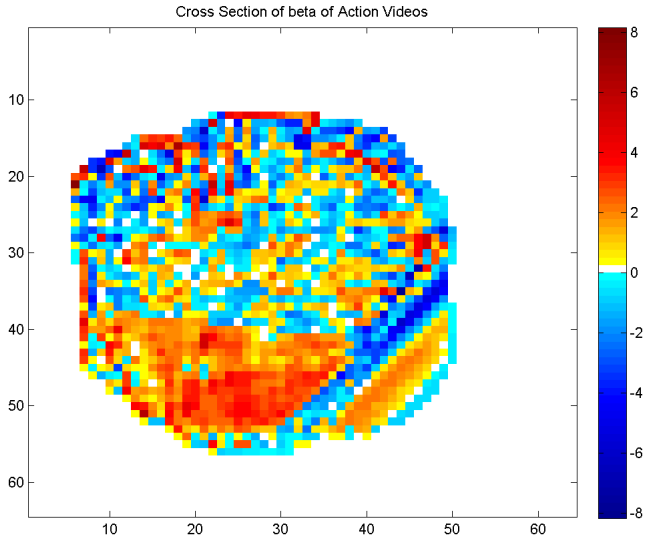

Note the bands of high and low correlation
at the lower right extremity.

We also ran into issues with processing speed and memory while performing analysis on these datasets. In order to obtain high resolution fMRI data, the dataset was extremely large and consumed large amounts of

memory. While we would have liked to process a higher quantity of information with more predictors for our linear model, constraints on time as well as physical computational limitations prevented us from doing so.

Future Work

In the short term, we want to extend the range of fitting functions for the generalized linear model. Currently, we are only using the convolution of the hemodynamic response function and the unit step, but by including more modeling terms, we could improve upon the accuracy of our detection methods.

In the future, we would like to extend this work into image reconstruction techniques for fMRI. By understanding the types of stimulation from different types of images and videos, we could hope to recreate the image that anyone is viewing at any given time from their fMRI data.

The Team

Ben Adler: Bioengineering and Biochemistry
Ben researched the background of fMRI function and the method of brain mapping via convolution.

Caitlin Makatura: Bioengineering
Caitlin researched the GLM model and convolution method, implemented the code for these sections, and wrote the background of the project.

Navin Pathak: Bioengineering
Navin implemented the code for organizing the data and collecting the results, analyzed the results, and uploaded the report to connexions.

Acknowledgements
We would like to thank Ricky Savjani and Mingbo Cai at the Baylor College of Medicine for providing us with the data for this project as well as guidance on approaching the problem.